


MIT LIBRARIES



3 9080 00672155 6





Digitized by the Internet Archive
in 2011 with funding from
Boston Library Consortium Member Libraries

<http://www.archive.org/details/existenceofsubga00harr>

HB31
.M415
no. 570

**working paper
department
of economics**

APR 1 1991

**THE EXISTENCE OF SUBGAME-PERFECT EQUILIBRIUM
IN GAMES WITH SIMULTANEOUS MOVES**

Christopher Harris

Number 570

December 1990

**massachusetts
institute of
technology**

**50 memorial drive
cambridge, mass. 02139**



THE EXISTENCE OF SUBGAME-PERFECT EQUILIBRIUM
IN GAMES WITH SIMULTANEOUS MOVES

Christopher Harris

Number 570

December 1990

M.I.T. LIBRARIES

APR 01 1991

RECEIVED

The Existence of Subgame-Perfect Equilibrium
in Games with Simultaneous Moves:
a Case for Extensive-Form Correlation

Christopher Harris^{*}

Nuffield College, Oxford

December 1990[†]

Abstract

The starting point of this paper is a simple, regular, dynamic game in which subgame-perfect equilibrium fails to exist. Examination of this example shows that existence would be restored if players were allowed to observe public signals. The main result of the paper shows that allowing the observation of public signals yields existence, not only in the example, but also in an extremely large class of dynamic games.

* I should like to thank Drew Fudenberg, Daniel Maldoom, Andreu Mas-Colell, Eric Maskin, Jean-Francois Mertens, Meg Meyer, Jean Tirole, John Vickers, and participants at seminars and workshops at Oxford, Stonybrook and Harvard for helpful comments on the earlier paper from which the present paper is drawn.

† The present paper is drawn from an earlier paper entitled "The Existence of Subgame-Perfect Equilibrium with and without Markov Strategies: a Case for Extensive Form Correlation".

1 Introduction

Consider the following example of a dynamic game. Firms set out with exogenously specified capacities (which are known to all). In period one they simultaneously choose investment levels (possibly on a random basis), and are then informed of one another's choices. The result is a change in capacities. In period two firms simultaneously choose production levels (within their capacity constraints), and are then informed of one another's choices. The result is a change in inventories. In period three firms simultaneously choose prices. The vector of prices chosen affects the vector of demands for their products, but so do certain exogenous random factors. Firms are informed of one another's chosen prices and of the final realised demands. (The demands bring about a second change in inventories.) The three period cycle of choices then begins afresh. And so on.

This game is an example of a game of the following general type. Time is discrete. There is a finite number of active players. There is also a passive player, Nature. In any given period: all players (both active and passive) know the outcomes of all previous periods; the set of actions available to any active player is compact, and depends continuously on the outcomes of the previous periods; the distribution of Nature's action (which is given exogenously) depends continuously on the outcomes of the previous periods; the players (active and passive) choose their actions simultaneously; and the outcome of the period is simply the vector of actions chosen. The outcome of the game as a whole is the (possibly infinite) sequence of outcomes of all periods, and players' payoffs are bounded and depend continuously on the outcome of the game.

This is as regular a class of dynamic games as one could ask for. A counter-example shows, however, that games in this class need not have a subgame-perfect equilibrium. It is therefore necessary to extend the equilibrium concept in such a way that existence is restored.¹

¹ The problem is reminiscent of that of the non-existence of Nash equilibrium in pure

What is the most natural extension of the equilibrium concept? One clue as to the answer to this question is provided by the following considerations. First, if we simplify the class of games under consideration by requiring that players' action sets are always finite, then subgame-perfect equilibrium always exists. Hence the non-existence problem appears to relate specifically to the fact that we allow a continuum of actions. Secondly, suppose instead that we simplify the class of games under consideration by assuming that players' action sets are independent of the outcomes of previous periods. Then a natural way of approximating a game is to consider subsets of players' action sets that consist of large but finite numbers of closely spaced actions.² Moreover, if one takes a sequence of increasingly fine approximations, and a subgame-perfect equilibrium of each of the approximations, then it is natural to expect that any limit point of the sequence of equilibrium paths so obtained will be an equilibrium path of the original game.³ One can, however, find examples in which the limit point involves a special form of extensive-form correlation: at the outset of each period agents observe a public signal.⁴

strategies. But there is one important difference: if agents' action sets are always finite, then subgame-perfect equilibrium does exist.

² Hellwig and Leininger (1986) were the first to introduce such an approximation, in the context of finite-horizon games of perfect information. They proved an upper-semicontinuity result: they showed that any limit point of equilibrium paths of the finite approximations is an equilibrium path of the original game. Their result can, however, be understood in terms of the results of Harris (1985a). Indeed, in that paper it is shown that the point-to-set mapping from period- t subgames of a game of perfect information to period- t equilibrium paths is upper semicontinuous. Hence, in order to obtain their result, one need only introduce a dummy player who chooses $n \in \mathbb{N} \cup \{\infty\}$ at the outset of the game. If $n < \infty$ then the remaining players will be restricted to actions chosen from the n^{th} approximation to their action sets. If $n = \infty$ then they will be free to choose actions from their original action sets.

³ Convergence of equilibrium paths, as used implicitly in Harris (1985a) and explicitly in Hellwig and Leininger (1986) and Börgers (1989) seems more relevant than the convergence of strategies considered by Fudenberg and Levine (1983) and Harris (1985b).

⁴ Fudenberg and Tirole (1985) consider two such games. They point out that the obvious discretisations of these games have a unique symmetric subgame-perfect equilibrium, and that the limiting equilibria obtained as the discretisation becomes arbitrarily fine involve correlation. Their games do not, however, yield counterexamples to the existence of subgame-perfect equilibrium. For both games possess asymmetric equilibria which do not involve correlation.

These considerations suggest that, at the very least, the equilibrium concept should be extended to allow the observation of public signals. This minimal extension is also sufficient. For, first, it restores existence. Secondly, with it, the equilibrium correspondence of a continuous family of games is upper semicontinuous. In particular, any limit point of equilibrium paths of finite approximations to a game whose action sets are continua is an equilibrium path of the game. Thirdly, any limit point of subgame-perfect ϵ -equilibrium paths of a game is an equilibrium path in the extended sense.⁵

The basic structure of the proof of the existence of subgame-perfect equilibrium in the extended sense is the same as the structure of the proof of the existence of subgame-perfect pure-strategy equilibrium in games of perfect information given in Harris (1985a). It breaks down into two main parts, the backwards and the forwards programs. The backwards program is most easily explained in the context of a game with finite horizon T . In such a game, it consists in solving recursively for what turn out to be the equilibrium paths of the game. More precisely, one solves first for the equilibrium paths of period- T subgames. (The set of equilibrium paths of a period- T subgame is the convex hull of the set of probability distributions over period- T outcomes.) Then one solves for the equilibrium paths of period- $(T-1)$ subgames. (The set of equilibrium paths of a period- $(T-1)$ subgame is a set of probability distributions over period- $(T-1)$ histories, i.e. a joint distribution over period- $(T-1)$ and period- T outcomes.) And so on until period 1 is reached, and a set of probability distributions over period-1 histories is obtained.

There is of course no general guarantee that the paths obtained in the course of carrying out the backwards program really are equilibrium paths until equilibrium strategies generating them have been constructed. To construct such strategies is the purpose of the forwards program. The essential idea is this. Pick any of the period-1 paths obtained as the final product of the backwards program. Such a path is a probability distribution over period-1 histories. Its marginal over period-1 outcomes can be used to

⁵ Chakrabarti (1988) claims that subgame-perfect ϵ -equilibria exist. We believe this claim, but are unable to vouch for the proof.

construct period–1 behaviour strategies for the players. Its conditional over period–2 histories yields continuation paths to be followed from period 2 onwards in the event that players do not deviate from their period–1 strategies. In order to obtain continuation paths in the event that some player does deviate from his period–1 strategy, it is sufficient to choose from among the period–2 paths obtained in the course of the backwards program some path which minimises the continuation payoff of that player. In this way continuation paths for all period–2 subgames are obtained. Period–2 strategies are then obtained from the marginals of these paths over period–2 outcomes. And so on until finally period–T strategies are obtained.

Implementing this proof does, however, involve two significant new complications. First, for technical reasons, it is necessary to maintain the induction that the point–to–set mapping from period–t subgames to period–t equilibrium paths is upper semicontinuous. This follows from a generalisation of the work of Simon and Zame (1990). Secondly, because players may randomise, it is essential to ensure that all the strategies constructed are measurable.⁶ That this requirement can be met emerges as a natural corollary of the construction employed: one can always construct a measurable family of conditional distributions from a measurable family of probability distributions.⁷

The organisation of the paper is as follows. Section 2 sets out the counterexample to the existence of subgame–perfect equilibrium. This example involves a three–period game with two players in each period. In this game, each player’s action set is compact

⁶ Hellwig and Leininger (1987) were the first to draw attention to the desirability of ensuring that strategies are measurable. Their approach to the existence of subgame–perfect equilibrium is not, however, necessary in order to obtain measurability. Indeed, if the assumptions of Harris (1985a) are specialised appropriately (by assuming that the embedding spaces for players’ action sets are compact metric instead of compact Hausdorff), then it is simple to show that exactly the construction given there can be carried out in a measurable way. Indeed, in the notation of Harris (1985a; p.625), all that is necessary is to ensure that the functions g_t^i can be chosen to be measurable; and that this is possible follows from the argument given in the first paragraph of the proof of Lemma 4.1 in the present paper.

⁷ Had we employed a dynamic–programming approach, the problem of ensuring measurability of strategies would have been much harder – see below for further discussion. Indeed, we do not know how to solve the measurable selection problem that arises when such an approach is employed.

and independent of the outcomes of previous periods, and each player's payoff function is continuous. Section 3 formulates the basic framework used in this paper for the analysis of dynamic games. Section 4 shows that when the dynamics are continuous, players' payoff functions are continuous, and the framework is extended to allow the observation of public signals, subgame-perfect equilibrium exists. This is the main result of the paper. Section 4.1 provides a detailed overview of the proof. Section 4.2 develops the analysis necessary for the backwards program. Section 4.3 develops the analysis necessary for the forwards program. Section 4.4 exploits the analysis of Sections 4.2 and 4.3 to prove the main result, showing in particular how to deal with the infinite-horizon case. Section 5 attempts to develop a perspective on the main result. It is shown there that there are at least three types of game in which the introduction of public signals is not necessary to obtain the existence of subgame-perfect equilibrium: games of perfect information, finite-action games, and zero-sum games. It is also shown that the introduction of public signals accommodates all the equilibria that one can obtain by taking finite-action approximations to a game, or by considering ϵ -equilibria of a game. We do not know whether the introduction of public signals is the minimal extension with these two properties, but it seems plausible to argue that it is.

2 The Counterexample

A counterexample to the existence of subgame–perfect equilibrium in the class of dynamic games described in the introduction must have a certain minimum complexity. First of all, it must have at least two periods. Otherwise the standard existence theorem for Nash equilibrium would apply. Secondly, there must be at least two players in the second period. For with only one player in the second period, the correspondence describing the continuation payoffs available to the player or players in period one would be convex valued. (The player in period two can randomise over any set of alternatives between which he is indifferent.) So the existence result of Simon and Zame (1990) would apply to period one. Thirdly, there must be at least two players in period one. For the continuation–payoff correspondence for period one will be upper semicontinuous, and a single player will therefore be able to find an optimum.

The counterexample presented in this section is not this minimal. It has three periods, with two players in each period. It would be of considerable interest to find an example with only two periods and two players in each period. For, aside from being minimal in relation to the class of dynamic games considered in this paper, such an example would settle the question as to whether equilibrium exists in two–stage games or not.

The cast of the counter–example, and their choice variables, are as follows. In period one two punters A and B pick $a \in [0,1]$ and $b \in [0,1]$ respectively. In period two, two greyhounds C and D each receive an injection of size $a + b$, which changes their attitude to the race. They pick $c \in [0,1]$ and $d \in [0,1]$, which are the times in which they complete the course. In period three each of two referees E and F must declare a winner, picking $e \in \{C,D\}$ and $f \in \{C,D\}$ respectively. In each period choices are made simultaneously, and players in later periods observe the actions taken by players in earlier periods.

Punter A obtains a payoff of $1 - a$ if greyhound C is declared the winner by both referees, and $-1 - a$ otherwise. Similarly, punter B obtains $1 - b$ if both referees declare D to be the winner, and $-1 - b$ otherwise. In other words, A wants C to win, B wants D to win, and both want a result. They would also like to keep their contributions to the injection as small as possible. The payoff to greyhound C is $2c$ if $e = C$ and $1 - (a + b)(1 - c)$ if $e = D$. That is, the form of his payoff depends on whether he or the other greyhound is declared the winner by referee E, but either way he would prefer to run the race as slowly as possible. Also, he would prefer to be first rather than second provided that he does not have to run too fast. The payoff to greyhound D is $2d$ if $e = D$ and $1 - (a + b)(1 - d)$ if $e = C$: like greyhound C, he is only interested in the verdict of referee E. Lastly, referee E gets payoff d if he declares C to be the winner and c if he declares D to be the winner. Referee F's payoffs are identical.

It will be seen that, in this game, each player's action set is compact and independent of the actions chosen by earlier players, and that players' payoffs are continuous.

The game centres around the two greyhounds. Because referee E chooses C when $c < d$ and D when $c > d$, the race between them is essentially a game of timing with the discontinuous payoffs $(2c, 1 - (a + b)(1 - d))$ if $c < d$, $(1 - (a + b)(1 - c), 2d)$ if $c > d$, and some convex combination of these two payoffs when $c = d$. (Note that the weights in the convex combination can depend on (c, d) .) Standard considerations therefore yield the following lemma.

Lemma 2.1 Suppose that $a + b > 0$. Then both C and D use mixed actions with distribution function G given by $G(x) = 0$ for $x \in [0, \frac{1}{2}]$ and $G(x) = (2x - 1)/(2x - 1 + (a + b)(1 - x))$ for $x \in [\frac{1}{2}, 1]$. \square

Notice that this mixed action is non-atomic, that its support is $[\frac{1}{2}, 1]$, and that all its probability mass concentrates at $\frac{1}{2}$ as $a + b \rightarrow 0+$. Not surprisingly, then, we obtain:

Lemma 2.2 Suppose that $a + b = 0$. Then both C and D choose $\frac{1}{2}$ with probability one. \square

The remainder of the argument is straightforward. If $a + b > 0$ then c will equal d with probability zero. Hence both referees will always agree. Also, each greyhound will win exactly half the time. The payoffs to A and B are therefore $-a$ and $-b$. If, on the other hand, $a + b = 0$, then $c = d = \frac{1}{2}$ with probability one. Hence each referee is indifferent as to which greyhound he declares to be the winner. Suppose that E opts for C with probability p and D with probability $1 - p$, and that F opts for C with probability q and D with probability $1 - q$. Then A's payoff is $2pq - 1$ and B's is $2(1 - p)(1 - q) - 1$.

It is easy to see, however, that the game with these payoffs between A and B has no equilibrium. Indeed, any $a > 0$ is strictly dominated for A, as is any $b > 0$ for B. So the only possibility for an equilibrium is for A to set $a = 0$ with probability one and B to set $b = 0$ with probability one. But if $b = 0$ then it must be that $2pq - 1 \geq 0$ — otherwise A would raise a from zero. Similarly, $2(1 - p)(1 - q) - 1 \geq 0$. And these two inequalities are mutually inconsistent.⁸ This contradiction establishes the counterexample.

The essence of the counterexample is this. As long as $a + b > 0$, both greyhounds use strictly mixed actions. This behaviour on their part generates a public signal endogenously within the game. The two referees use this public signal to co-ordinate their actions. When $a + b = 0$, however, this signal suddenly degenerates, and the only way in which the referees can coordinate their actions is by both choosing C with probability one or both choosing D with probability one. But if they both choose C then B gets

⁸ To see this, note that the sum of the payoffs of the two players A and B is 0 if $e = f$ and -1 if $e \neq f$. Hence the expected value of this sum is non-positive. If the expected value is strictly negative, then the payoff of at least one player is also strictly negative. If the expected value of the sum is zero, then e and f must be perfectly correlated. In this case the payoff of the player against whom the decision goes is -1 .

-1 . So he would prefer the truly random outcome. Similarly, if both referees choose D then A will wish to restore the truly random outcome.

In the light of this explanation, it is natural to try to restore existence by allowing players to observe suitable public signals. In the present example, this would amount to allowing the two referees to toss a coin to determine the winner in the event of a tie (which would certainly restore existence). That such an extension yields existence in the general case will be demonstrated in Section 4.

3 The Basic Framework

3.1 The Data

There is a non-empty finite set \mathcal{A} of active players, indexed by i or j . There is also a passive player, Nature, whose index is i or $j = \emptyset$. The overall set of players is therefore $\mathcal{I} = \mathcal{A} \cup \{\emptyset\}$. Time is discrete, and is indexed by t or $s \in \mathcal{T} = \{0, 1, 2, \dots\}$. For notational reasons, we assume that $\mathcal{I} \cap \mathcal{T} = \emptyset$. The players play one of a family of games, parameterised by $x^0 \in X^0$. In each active period $t \in \mathcal{A} = \{1, 2, \dots\}$ each player $i \in \mathcal{I}$ must choose an action. The set of actions available to her or him depends on which game is being played, and on the outcomes of previous periods. This situation is modelled by a point-to-set mapping $A_{ti}: X^{t-1} \rightarrow Y_{ti}$. The vector $x^{t-1} \in X^{t-1}$ lists the parameter of the game being played, and the outcomes of any preceding periods, while $A_{ti}(x^{t-1}) \subset Y_{ti}$ is the set of actions available. For reasons of expositional economy, we take it that $A_{t\emptyset}(x^{t-1}) = Y_{t\emptyset}$ for all $x^{t-1} \in X^{t-1}$.⁹ The set of outcomes possible in period t is nothing more than the set of profiles of actions that players can take. It is modelled by the point-to-set mapping $A_t: X^{t-1} \rightarrow Y_t$, where $Y_t = \times_{i \in \mathcal{I}} Y_{ti}$ and $A_t(x^{t-1}) = \times_{i \in \mathcal{I}} A_{ti}(x^{t-1})$ for all $x^{t-1} \in X^{t-1}$. And X^t is simply the set of all pairs (x^{t-1}, y_t) such that $x^{t-1} \in X^{t-1}$ and $y_t \in A_t(x^{t-1})$, i.e. the graph of A_t . Finally, the payoff of any active player i depends on which game $x^0 \in X^0$ is played, and the outcomes of all periods. It is modelled by a function $u_i: X^\omega \rightarrow \mathbb{R}$. Here X^ω is the set of vectors $x = (x_0, x_1, \dots) \in X^0 \times \left[\times_{t=1}^\omega Y_t \right]$ such that $x^{t-1} = (x_0, x_1, \dots, x_{t-1}) \in X^{t-1}$ for all $t \geq 1$. (Nature does not have a payoff function.)

We make the following standing assumptions about these data:

- (i) X^0 , and all of the sets Y_{ti} , are non-empty complete separable metric spaces;
- (ii) for all $i \in \mathcal{A}$ and all $t \geq 1$, $A_{ti}: X^{t-1} \rightarrow \mathcal{K}(Y_{ti})$ is measurable;
- (iii) for all $i \in \mathcal{A}$, u_i is bounded and measurable.

⁹ This assumption does in principle involve a slight loss of generality. However, since we shall in any case fix Nature's strategy below, there seems to be little advantage in constraining her actions as well.

(In this paper, ‘measurable’ will always mean Borel measurable unless explicitly stated to the contrary; and $\mathcal{K}(\cdot)$ will always denote the space of non-empty compact subsets of \cdot endowed with the Hausdorff metric.) These assumptions are designed to ensure that every player possesses at least one strategy, and that associated with every strategy profile and every subgame there is a well defined payoff for every active player (see below). Stronger assumptions are needed to ensure the existence of equilibrium.

3.2 Strategies and extended strategies

In the standard version of our model, the evolution of the game is as follows. At the beginning of period 1, players are told which game they are playing. That is, they are informed of x^0 . They then simultaneously choose actions from their action sets for period 1. This completes period 1. At the beginning of period 2, players are told the outcome of period 1. That is, they are told y_1 . They then choose simultaneously from their action sets for period 2. This completes period 2. And so it goes on. It follows that players’ information in period t can be summarised by a vector $x^{t-1} \in X^{t-1}$, and that X^{t-1} can be identified with the set of subgames in period t .

Definition 3.1 For each $t \geq 1$ and each $i \in \mathcal{I}$, a strategy for player i in period t is a measurable function $f_{ti}: X^{t-1} \rightarrow \mathcal{PM}(Y_{ti})$ such that $\text{supp}[f_{ti}(x^{t-1})] \subset A_{ti}(x^{t-1})$ for all $x^{t-1} \in X^{t-1}$.

Here $\mathcal{PM}(Y_{ti})$ denotes the set of probability measures over Y_{ti} . (In this paper, ‘probability measure’ will always mean Borel probability measure; and spaces of probability measures will be taken to be endowed with the weak topology¹⁰ unless explicitly stated to the contrary.)

¹⁰ In the terminology of Parthasarathy (1967), a sequence of probability measures $\{\lambda^n\} \subset \mathcal{PM}(Y_t)$ converges in the weak topology to λ iff $\int \chi d\lambda_n$ converges to $\int \chi d\lambda$ for all continuous bounded $\chi: Y_t \rightarrow \mathbb{R}$.

The f_{ti} are behaviour strategies. For each t , i and x^{t-1} , $f_{ti}(x^{t-1})$ can be thought of as the randomising device that player i will use to choose among the actions available to him in period t when x^{t-1} is the previous history of the game. As such, it is independent of the devices used by the other players in period t , and of all devices used in all preceding and subsequent periods.

Strategies and strategy profiles certainly exist under our standing assumptions. Moreover, given any strategy profile and any subgame, the payoffs of active players can be calculated in the natural way. We may therefore define subgame–perfect equilibrium as follows.

Definition 3.2 A subgame–perfect equilibrium is a strategy profile $\langle f_{ti} | t \geq 1, i \in \mathcal{I} \rangle$ such that, for all $t \geq 1$, all $x^{t-1} \in X^{t-1}$, and all $i \in \mathcal{I}$, player i cannot improve his payoff in subgame x^{t-1} by a unilateral change in his strategy.

As the counterexample of Section 2 shows, however, subgame–perfect equilibrium need not exist in general. We therefore introduce the concept of an extended strategy. Suppose that $\{\xi_t | t \geq 1\}$ is a sequence of signals, drawn independently from $[0,1]$ according to the uniform distribution. Suppose that, at the beginning of each active period $t \geq 1$, players are informed not only of the outcome of the preceding period, but also of ξ_t . Then the information available to them when they make their choices can be summarised by a vector $h^{t-1} = (x^{t-1}, \xi^t)$, where $\xi^t = (\xi_1, \dots, \xi_t)$. We refer to the set $H^{t-1} = X^{t-1} \times [0,1]^t$ of such vectors as the set of extended subgames of the game in period t .

Definition 3.3 For each $t \geq 1$ and each $i \in \mathcal{I}$, an extended strategy for player i in period t is a measurable function $f_{ti}: H^{t-1} \rightarrow \mathcal{PM}(Y_{ti})$ such that $\text{supp}[f_{ti}(x^{t-1}, \xi^t)] \subset A_{ti}(x^{t-1})$ for all $(x^{t-1}, \xi^t) \in H^{t-1}$.

Once again the f_{ti} are behaviour strategies. So $f_{ti}(x^{t-1}, \xi^t)$ can be thought of as the randomising device that player i will use to choose among the actions available to him in period t , when x^{t-1} is the previous history of the game and ξ^t is the vector of public signals observed up to and including period t . But this time this device is independent, not only of the devices used by the other players in period t and of all devices used in all other periods, but also of all the public signals.

Definition 3.4 A subgame–perfect equilibrium in extended strategies is an extended–strategy profile $\langle f_{ti} | t \geq 1, i \in \mathcal{N} \rangle$ such that, for all $t \geq 1$, all $h^{t-1} \in H^{t-1}$, and all $i \in \mathcal{N}$, player i cannot improve his payoff in extended subgame h^{t-1} by a unilateral change in his strategy.

A subgame–perfect equilibrium in extended strategies allows a limited amount of coordination: after the outcome of a period has been realised, the players can coordinate on the continuation equilibrium to be played following that outcome by exploiting the next public signal. It is therefore a limited kind of extensive–form correlated equilibrium. (The correlation is limited in that the players observe a common signal, rather than each privately observing a single component of a vector of signals.)

4 The Main Result

We need the following assumptions:

- (A1) for all $t \geq 1$ and all $i \in \mathcal{I}$, the mapping $A_{ti}: X^{t-1} \rightarrow \mathcal{K}(T_{ti})$ is continuous;
- (A2) for all $t \geq 1$, a strategy $f_{t\emptyset}$ (and not an extended strategy) for Nature in period t is given, and $f_{t\emptyset}: X^{t-1} \rightarrow \mathcal{PM}(Y_{t\emptyset})$ is continuous;
- (A3) for all $i \in \mathcal{I}$, u_i is bounded and continuous.

Taken in conjunction with the standing assumptions made in respect of the basic framework, they guarantee the existence of a subgame–perfect equilibrium in extended strategies.

The present section is devoted to a proof of this fact. It begins, in Section 4.1, with an overview of the proof. The proof involves two main steps. The analysis necessary for the first of these, namely the backwards program, is given in Section 4.2. The analysis necessary for the other, namely the forwards program, is given in Section 4.3. Section 4.4 then integrates the backwards and the forwards programs to complete the proof. The reader may wish to read Section 4.1 and then skip to Section 5.

4.1 Overview

Imagine for the purposes of the present subsection that the game has horizon T . Then the proof divides into essentially two steps. In the first step, one programs backwards from stage T , finding what turn out to be the equilibrium paths of the game. In the second, one programs forwards, constructing strategies that generate these paths and that constitute an equilibrium.

The basic ideas of the backwards program are as follows. First, for each history $x^{T-1} \in X^{T-1}$, find the set of Nash equilibria (including Nash equilibria in mixed strategies) that can occur in the final stage of the game. Let $C_T(x^{T-1})$ be the set of probability distributions over Y_T that result from these Nash equilibria. Because action sets are compact and payoff functions are continuous, $C_T(x^{T-1})$ is a non–empty compact

set contained in $\mathcal{PM}(Y_T)$. Because actions sets depend continuously on x^{T-1} , the point-to-set mapping C_T from X^{T-1} to $\mathcal{K}(\mathcal{PM}(Y_T))$ is upper semicontinuous.¹¹

Next, let x^{T-2} be the history of the game prior to period $T-1$. For any $y_{T-1} \in A_{T-1}(x^{T-2})$, the set of continuation paths consistent with equilibrium is given by $\text{co}[C_T(x^{T-2}, y_{T-1})]$. (This set consists, in effect, of those probability distributions over Y_T that can be obtained by randomising over Nash equilibria in $C_T(x_{T-1})$.) For each selection $c_T(x^{T-2}, \cdot)$ from the correspondence $\text{co}[C_T(x^{T-2}, \cdot)]$, find the set of Nash equilibria that can occur in period $T-1$ when the continuation of the game is specified by $c_T(x^{T-2}, \cdot)$. (For a general $c_T(x^{T-2}, \cdot)$, this set may be empty. But there will always be at least one $c_T(x^{T-2}, \cdot)$ for which it is not. This is essentially the content of the existence theorem of Simon and Zame (1990).) And for each Nash equilibrium associated with $c_T(x^{T-2}, \cdot)$, find the probability distribution over $Y_{T-1} \times Y_T$ that results when that equilibrium is played in period $T-1$ and the continuation paths are given by $c_T(x^{T-2}, \cdot)$. Let $C_{T-1}(x^{T-2})$ be the set of such distributions obtained when both $c_T(x^{T-2}, \cdot)$, and the Nash equilibrium associated with it, are allowed to vary. Then $C_{T-1}(x^{T-2})$ is a non-empty compact subset of $\mathcal{PM}(Y_{T-1} \times Y_T)$, and the point-to-set mapping C_{T-1} from X^{T-2} to $\mathcal{K}(\mathcal{PM}(Y_{T-1} \times Y_T))$ is upper semicontinuous. (This is essentially the content of the generalisation of Simon and Zame's theorem given in Section 4.2.)

¹¹ The relevance of the assumption that players' action sets in period T depend continuously on the history of the game prior to period T emerges clearly in the context of a two-period game. For each $n \geq 1$, let a_2^n be a Nash equilibrium in period 2 of such a game when the outcome of period 1 is y_1^n , and suppose that $(y_1^n, a_2^n) \rightarrow (y_1, a_2)$ as $n \rightarrow \infty$. Because player i 's action set is upper semicontinuous in the outcome of period 1, the action a_{2i} is actually available to him in period 2 when the outcome of period 1 is y_1 . Also, because his action set in period 2 is lower semicontinuous in the outcome of period 1, any deviation open to him in period 2 when the outcome of period 1 is y_1 can be approximated by deviations open to him when the outcome of period 1 is y_1^n . Hence a_{2i} really is an optimal action.

Finally, iterate until period 1 is reached, and a set $C_1(x^0)$ of probability distributions over $\times_{t=1}^T Y_t$ is obtained for all $x^0 \in X^0$. (The set $\text{co}[C_1(x^0)]$ will turn out to be the set of equilibrium paths of the game x^0 .) This completes the backwards program.

Consider now the forwards program. Suppose that, for each $x^0 \in X^0$, we are given $c_1(x^0) \in \text{co}[C_1(x^0)]$. Since any path in $\text{co}[C_1(x^0)]$ can be viewed as a randomisation over paths in $C_1(x^0)$, we can associate a probability measure $d_1(x^0)$ over $C_1(x^0)$ with $c_1(x^0)$. Also, for each $d_1(x^0)$ we can find a random variable $e_1(x^0, \cdot): [0,1] \rightarrow C_1(x^0)$ whose distribution over $C_1(x^0)$ is $d_1(x^0)$ when $[0,1]$ is given Lebesgue measure. For each i and each $\xi^1 \in [0,1]$, let $f_{1i}(x^0, \xi^1)$ be the marginal of $e_1(x^0, \xi^1)$ over Y_{1i} . By definition of $C_1(x^0)$, there exists a selection $c_2(x^0, \cdot, \xi^1)$ from $\text{co}[C_2(x^0, \cdot)]$ such that: (i) the $f_{1i}(x^0, \xi^1)$ constitute a Nash equilibrium in period 1 when the continuation paths are given by $c_2(x^0, \cdot, \xi^1)$; and (ii) for all y_1 , $c_2(x^0, y_1, \xi^1)$ is the conditional distribution of $e_1(x^0, \xi^1)$ given that y_1 is the outcome in period 1.

Similarly, beginning with $c_2(x^1, \xi^1) \in \text{co}[C_2(x^1)]$ one obtains $d_2(x^1, \xi^1)$, $e_2(x^1, \xi^2)$, $f_{2i}(x^1, \xi^2)$ and $c_3(x^2, \xi^2)$. And so on until one has a complete set of strategies f_{ti} for all $1 \leq t \leq T$ and $i \in \mathcal{N}$. That these strategies really do produce the paths they are intended to produce is primarily a technical matter. The essential points are that the expectation of $e_t(x^{t-1}, \xi^{t-1}, \cdot)$ is always $c_t(x^{t-1}, \xi^{t-1})$, and that $c_{t+1}(x^{t-1}, y_t, \xi^t)$ is always the conditional distribution of $e_t(x^{t-1}, \xi^t)$ given y_t . That they constitute an equilibrium follows from the fact that the $f_{ti}(x^{t-1}, \xi^t)$ constitute a Nash equilibrium. This ensures that no single-period deviation is profitable. That no more complex deviation is profitable follows from a standard unravelling argument. This completes the forwards program.

Overall, the essential ingredients of the theorem are the following: players action sets are compact; players' payoffs are continuous; Nature's strategy is continuous; players' action sets in a given period depend continuously on the history of the game prior to that period; players observe public signals; and players' payoffs are bounded. The role played

by the first three of these should be obvious. The role played by the fourth is, roughly, to ensure that the point-to-set mapping from subgames in period t to equilibrium paths of those subgames is upper semicontinuous. The role of the fifth is to ensure that the point-to-set mapping from subgames in period t to equilibrium paths of those subgames is, in addition, convex valued. And the sixth is needed only for the relatively technical reason that we do not insist that Nature's behaviour strategies have compact support.

4.2 The backwards program

In this subsection we develop the analysis necessary for the backwards program. To this end:

- (i) let X , Y_i for all $i \in \mathcal{I}$, and Z be non-empty complete separable metric spaces;
- (ii) for all $i \in \mathcal{I}$, let $A_i: X \rightarrow \mathcal{K}(Y_i)$ be a continuous point-to-set mapping;
- (iii) let $f_\emptyset: X \rightarrow \mathcal{PM}(Y_\emptyset)$ be a continuous mapping;
- (iv) let $Y = \times_{i \in \mathcal{I}} Y_i$ and let $A(x) = Y_\emptyset \times \left[\times_{i \in \mathcal{I}} A_i(x) \right]$ for all $x \in X$;
- (v) let $C: \text{gr}(A) \rightarrow \mathcal{K}(Z)$ be an upper semicontinuous point-to-set mapping;
- (vi) for each $i \in \mathcal{I}$, let $u_i: \text{gr}(C) \rightarrow \mathbb{R}$ be a continuous function.

The similarity of this notation to that of Section 3 is deliberate, and is intended to be helpful, not confusing.

With reference to the model of Section 3 and the discussion of Section 4: X is the set of past histories; A_i describes the dependence of player i 's action set on the past history; f_\emptyset describes Nature's behaviour; C describes the set of possible continuation paths; and u_i is the payoff of active player i . In terms of a somewhat more abstract perspective, what we have is a continuous family of games parameterised by $x \in X$. The three essential ingredients of such a family are: the continuity of the action sets; the upper semicontinuity of the continuation paths; and the continuity of the payoffs in the parameter, the outcome and the continuation.

For each $x \in X$ we are interested in the set of equilibrium paths $\hat{\Psi}C(x) \subset \mathcal{PM}(Y \times Z)$ that are obtained when one is free to choose any randomisation over

$C(x,y)$ as the continuation path following the outcome $y \in A(x)$. More formally, $\lambda \in \widehat{\Psi}C(x)$ iff:

- (i) the marginal μ of λ over Y is a product measure with $\text{supp}[\mu] \subset A(x)$;
- (ii) the marginal μ_\emptyset of μ over Y_\emptyset is $f_\emptyset(\cdot|x)$;
- (iii) λ possesses an r.c.p.d. (regular conditional probability distribution) ν over Z such that $\text{supp}[\nu(\cdot|y)] \subset C(x,y)$ for all $y \in A(x)$;
- (iv) moreover the marginals μ_i of μ over the Y_i constitute a Nash equilibrium when the continuation path following outcome y is $\nu(\cdot|y)$ for all y .

(It may be helpful to state more explicitly what is meant by (iv). Suppose we are given probability measures μ_i over $A_i(x)$ for all $i \in \mathcal{I}$ and a transition probability ν such that $\text{supp}[\nu(\cdot|y)] \subset C(x,y)$ for all $y \in A(x)$. Let λ be the measure over $Y \times Z$ obtained by combining the product of the μ_i with ν . Then the payoff of an active player i is $\int u_i(x,y,z)d\lambda(y,z)$, and the μ_i constitute a Nash equilibrium if no such player can improve his payoff by changing his probability measure.)

Lemma 4.1 $\widehat{\Psi}C$ is an upper semicontinuous map from X into $\mathcal{K}(\mathcal{PM}(Y \times Z))$.

Proof We begin with some notation. For each i , x , and $y \in A(x)$, let $p_i(x,y) = \min\{u_i(x,y,z) | z \in C(x,y)\}$. Then $p_i(x,y)$ is the lowest payoff that can be imposed on player i in game x following outcome y . Standard considerations show that p_i is lower semicontinuous. Let $P_i(x,y) = \{z | z \in C(x,y) \text{ and } u_i(x,y,z) = p_i(x,y)\}$. We shall need a measurable selection from P_i . Following Deimling (1985; Proposition 24.3 and Theorem 24.3), the existence of such a selection will follow if $P_i^{-1}(D) = \{(x,y) | D \cap P_i(x,y) \neq \emptyset\}$ is measurable for all closed $D \subset Z$. To this end, define $p_i^D(x,y) = \min\{u_i(x,y,z) | z \in D \cap C(x,y)\}$. (By definition $p_i^D(x,y) = \infty$ if $D \cap C(x,y) = \emptyset$.) Like p_i , p_i^D is lower semicontinuous. Hence $P_i^{-1}(D) = \{(x,y) | p_i(x,y) = p_i^D(x,y)\}$ is measurable. Let $q_i: \text{gr}(A) \rightarrow Z$ be the resulting measurable selection.

The proof now divides into three parts. For each $x \in X$, let $B(x)$ be the set of paths of game x that are consistent with Nature's mixed action $f_\emptyset(\cdot|x)$. The first part of the proof shows that this mapping is an upper semicontinuous mapping from X into $\mathcal{H}(\mathcal{PM}(Y \times Z))$. The second shows that the graph of $\hat{\Psi}C$ is closed. Since this graph is clearly contained in that of B , all that remains is to demonstrate that $\hat{\Psi}C$ is non-empty valued. This is done in the third step. The second step accounts for the bulk of the proof.

It is clear that B is non-empty valued, and that its graph is closed. All that we need therefore show is that it is compact valued. To this end, let $x \in X$ be given. Since $B(x)$ is closed, we need only show that it is tight in order to show that it is compact. That is, we need only show that, for all $\epsilon > 0$, there exists a compact $K \subset Y \times Z$ such that $\lambda(K) \geq 1 - \epsilon$ for all $\lambda \in B(x)$. But there certainly exists a compact $\tilde{K} \subset Y_\emptyset$ such that $f_\emptyset(\tilde{K}) \geq 1 - \epsilon$, and we may therefore let K be the graph of the restriction of $C(x, \cdot)$ to $\tilde{K} \times \left[\times_{i \in \mathcal{I}} A_i(x) \right]$. This completes the first step.

In order to begin the second step, suppose that $(x_n, \lambda_n) \in \text{gr}[\hat{\Psi}C]$, and that $(x_n, \lambda_n) \rightarrow (x, \lambda)$. For all i , let π_{ni} be player i 's payoff from λ_n , and let π_i be his payoff from λ . From the definition of $\hat{\Psi}C(x_n)$ one can deduce that:

$$\int \chi_i(y_i) \chi_j(y_j) d\lambda_n(y, z) = \left[\int \chi_i(y_i) d\lambda_n(y, z) \right] \left[\int \chi_j(y_j) d\lambda_n(y, z) \right] \quad \dots(5.1)$$

for all $i \neq j$ and all continuous $\chi_i : Y_i \rightarrow \mathbb{R}$ and $\chi_j : Y_j \rightarrow \mathbb{R}$;

$$\int \chi_\emptyset(y_\emptyset) d\lambda_n(y, z) = \int \chi_\emptyset(y_\emptyset) df_\emptyset(y_\emptyset | x_n) \quad \dots(5.2)$$

for all continuous $\chi_\emptyset : Y_\emptyset \rightarrow \mathbb{R}$;

$$\text{supp}(\lambda_n) \subset \text{gr}(C(x_n, \cdot)) \quad \dots(5.3),$$

where $C(x_n, \cdot)$ is regarded as a (possibly empty-valued) correspondence from Y to Z ($C(x_n, y)$ is empty iff $y \notin A(x_n)$);

$$\int \chi_i(y_i) u_i(x, y, z) d\lambda_n(y, z) = \pi_{ni} \int \chi_i(y_i) d\lambda_n(y, z) \quad \dots(5.4)$$

for all active i and all continuous $\chi_i: Y_i \rightarrow \mathbb{R}$; and

$$\pi_{ni} \geq \int p_i(x_n, y \setminus a_{ni}) d\lambda_n(y, z) \quad \dots(5.5)$$

for all active i and all $a_{ni} \in A_i(x_n)$.

Equation (5.1) follows from the independence of the marginals of λ_n over the Y_i . Equation (5.2) follows from the fact that λ_n is consistent with Nature's strategy. Relation (5.3) follows from the facts that the marginal μ_n of λ_n is supported on $A(x_n)$, and that its conditional $\nu_n(\cdot | y)$ is supported on $C(x_n, y)$ for all $y \in A(x_n)$. Equation (5.4) holds because the expectation of $u_i(x_n, \cdot, \cdot)$ conditional on y_i is π_{ni} a.s. This follows from the fact that, in equilibrium, the probability mass of player i 's mixed strategy must be confined to a set of actions, each of which yields him his equilibrium payoff. (It is important to avoid the word 'support' here. For player i 's continuation payoff $\int u_i(x_n, y, z) d\nu_n(z | y)$ need not be continuous in y , and so the set of pure strategies yielding him π_{ni} need not be closed.) Inequality (5.5) states that player i 's equilibrium payoff must be at least what he would get if he were to deviate to a_{ni} and if, following this deviation and no matter what the realised choices of the other players, the worst conceivable continuation path from his point of view occurred. Equation (5.4) captures the idea that no redistribution of probability mass among equilibrium actions is worthwhile. Inequality (5.5) captures the idea that no redistribution of probability mass from equilibrium to out-of-equilibrium actions is worthwhile.

Now it is easy to see that relations (5.1), (5.2) and (5.3) are preserved in the limit. Also, since $\pi_{ni} = \int u_i(x_n, y, z) d\lambda_n(y, z)$, $\pi_{ni} \rightarrow \pi_i$. Hence (5.4) is preserved too. Lastly, for

any $a_i \in A_i(x)$, there exist $a_{ni} \in A_i(x_n)$ such that $a_{ni} \rightarrow a_i$, by the lower semicontinuity of A_i . Since also w_i is lower semicontinuous, (5.5) too is preserved. The proof that $\widehat{\Psi}C$ has a closed graph will therefore be complete if we can show that the analogues of (5.1)–(5.5) for x and λ , which we shall refer to as (5.1')–(5.5'), imply that $\lambda \in \widehat{\Psi}C(x)$.

Using standard considerations we can deduce from (5.1'), (5.2') and (5.3') that the marginal μ of λ over Y is a product measure supported on $A(x)$, that the marginal of μ over Y_\emptyset is $f_\emptyset(\cdot|x)$, and that λ has an r.c.p.d. $\tilde{\nu}$ over Z such that $\tilde{\nu}(\cdot|y)$ is supported on $C(x,y)$ for all $y \in A(x)$. Now consider the game with continuation paths specified by $\tilde{\nu}$. The marginals μ_i of μ over the Y_i need not constitute a Nash equilibrium for this game. However, by (5.4'), the set \tilde{Y}_i of $a_i \in A_i(x)$ such that deviation by active player i from μ_i to δ_{a_i} , the Dirac measure concentrated at a_i , will raise his payoff above π_i has μ_i -measure zero. We may therefore alter $\tilde{\nu}$ in such a way that it remains an r.c.p.d. of λ , but such deviations are no longer attractive. We set $\nu(\cdot|y) = \tilde{\nu}(\cdot|y)$ if $y_i \in Y_i \setminus \tilde{Y}_i$ for all i , and $\nu(\cdot|y) = \delta_{q_i(x,y)}$ if $y_j \in Y_j \setminus \tilde{Y}_j$ for all $j \neq i$ but $y_i \in \tilde{Y}_i$. It is of no importance how we define $\nu(\cdot|y)$ if $y_i \in \tilde{Y}_i$ for more than one i . (Such a y corresponds to a coordinated deviation by two or more players.) For definiteness, we set $\nu(\cdot|y)$ equal to $\tilde{\nu}(\cdot|y)$ in this case. Inequality (5.5') ensures that, with the new continuation paths specified by ν , deviation is no longer profitable. The μ_i therefore constitute a Nash equilibrium with these continuations, and therefore $\lambda \in \widehat{\Psi}C(x)$.

Having proved that $\widehat{\Psi}C$ has closed graph, all that remains is to show that $\widehat{\Psi}C$ is non-empty valued. To this end, fix x and find, for each i , a sequence $\{y_{ni}\}$ that is dense in $A_i(x)$. For each $1 \leq N < \infty$, construct a finite game with action sets $\{y_{ni} | 1 \leq n \leq N\}$ in the natural way from the existing game x . And for $N = \infty$ simply take the existing game x . The family of games obtained as N varies is continuous, and its equilibrium correspondence therefore has compact values and a closed graph. Moreover standard results show that the set of equilibria is non-empty when $N < \infty$. It follows that

it is non-empty for $N = \infty$ too. \square

Lemma 4.1 captures the essence of the backwards program: we begin with an upper semicontinuous correspondence C depending on the sequence of outcomes up to and including the current period, and end up with an upper semicontinuous correspondence $\hat{\Psi}C$ depending on the sequence of outcomes preceding the current period. There is however one further complication. When we come to complete the proof of the theorem in Section 4.4, the continuation paths in Z will themselves be probability measures over sequences of future outcomes. In the present, more abstract, setting this can be captured by writing $Z = \mathcal{PM}(W)$, where W is itself a complete separable metric space. With this notation, the equilibrium paths described by $\hat{\Psi}C$ lie in $\mathcal{PM}(Y \times \mathcal{PM}(W))$. This means that $\hat{\Psi}C$ cannot serve directly as input for the next iteration of the backwards program. What is required instead is a correspondence $\Psi C : X \rightarrow \mathcal{PM}(Y \times W)$. (The point is that elements of $\mathcal{PM}(Y \times W)$ are probability measures over sequences of future outcomes, whereas those of $\mathcal{PM}(Y \times \mathcal{PM}(W))$ are not.) Such a correspondence can be obtained by replacing elements of $\mathcal{PM}(\mathcal{PM}(W))$ with their expectations.

More formally, for each $x \in X$ let $\Psi C(x) \subset \mathcal{PM}(Y \times W)$ be the set of measures λ such that:

- (i) the marginal μ of λ over Y is a product measure with $\text{supp}(\mu) \subset A(x)$;
- (ii) the marginal μ_\emptyset of μ over Y_\emptyset is $f_\emptyset(x)$;
- (iii) λ possesses an r.c.p.d. ν over W such that $\nu(\cdot | y) \in \text{co}[C(x,y)]$ for all $y \in A(x)$;
- (iv) moreover the marginals μ_i of μ over the Y_i constitute a Nash equilibrium when the continuation path following outcome y is $\nu(\cdot | y)$ for all y .

(Once again it may be helpful to expand on (iv). Suppose that we are given $\mu_i \in \mathcal{PM}(A_i(x))$ and a transition probability ν such that $\nu(\cdot | y) \in \text{co}[C(x,y)]$ for all $y \in A(x)$. Then active player i 's payoff is $\int u_i(x,y,\nu(\cdot | y))d\mu(y)$, where μ is the product of the μ_i . The μ_i constitute a Nash equilibrium if no such player can improve his payoff

by changing his μ_1 .)

Lemma 4.2 Suppose that the domain of definition of $u_1(x,y,z)$ is convex in z , and suppose that $u_1(x,y,z)$ is linear in z on its domain. Then ΨC is an upper semicontinuous mapping from X into $\mathcal{K}(\mathcal{PM}(Y \times W))$.

It should be noted that, when we come to apply Lemma 4.2, $u_1(x,y,z)$ will be defined as the expectation with respect to z of an underlying function of (x,y,w) . So the convexity of the domain of $u_1(x,y,z)$ in z , and the linearity of $u_1(x,y,z)$ in z , will be natural consequences of the problem structure. Lemma 4.2 will be proved by showing, roughly speaking, that ΨC is the image of $\hat{\Psi} C$ under the projection that consists of replacing elements of $\mathcal{PM}(\mathcal{PM}(W))$ by their expectations, and that this projection is continuous.

Proof Let $\lambda \in \mathcal{PM}(Y \times Z)$, let μ be the marginal of λ over Y , and let ν be an r.c.p.d of μ over Z . Let $\bar{\mu} = \mu$, and let $\bar{\nu}(\cdot | y) = \int z d\nu(z|y)$ be the expectation of $\nu(\cdot | y)$ for all $y \in Y$. Then the projection $\bar{\lambda}$ of λ onto $\mathcal{PM}(Y \times W)$ is the probability measure with marginal $\bar{\mu}$ and conditional $\bar{\nu}$. (It should be clear that $\bar{\lambda}$ is independent of the r.c.p.d. chosen for λ .) The first step of the proof is to show that projection is continuous.

To this end, let $\{\lambda_n\} \subset \mathcal{PM}(Y \times Z)$ converge to $\lambda \in \mathcal{PM}(Y \times Z)$. Then, to show that $\{\bar{\lambda}_n\} \subset \mathcal{PM}(Y \times W)$ converges to $\bar{\lambda} \in \mathcal{PM}(Y \times W)$, it suffices to show that

$$\int \chi(y)\psi(w) d\bar{\lambda}_n(y,w) \rightarrow \int \chi(y)\psi(w) d\bar{\lambda}(y,w)$$

for all continuous functions $\chi: Y \rightarrow \mathbb{R}$ and $\psi: W \rightarrow \mathbb{R}$. Let $\ell: \mathcal{PM}(W) \rightarrow \mathbb{R}$ be the continuous linear functional defined by $\ell(\kappa) = \int \psi(w) d\kappa(w)$. Then

$$\int \chi(y) \psi(w) d\bar{\lambda}_{\mathbf{n}}(y, w) = \int \left[\int \psi(w) d\bar{\nu}_{\mathbf{n}}(w|y) \right] \chi(y) d\bar{\mu}_{\mathbf{n}}(y)$$

(by Fubini's theorem)

$$= \int \left[\ell(\bar{\nu}_{\mathbf{n}}(\cdot|y)) \right] \chi(y) d\bar{\mu}_{\mathbf{n}}(y)$$

(by definition of ℓ)

$$= \int \left[\int \ell(z) d\nu_{\mathbf{n}}(z|y) \right] \chi(y) d\bar{\mu}_{\mathbf{n}}(y)$$

(by definition of the expectation of $\nu_{\mathbf{n}}(\cdot|y)$)

$$= \int \chi(y) \ell(z) d\lambda_{\mathbf{n}}(y, z)$$

(by Fubini again, and because $\mu_{\mathbf{n}} = \bar{\mu}_{\mathbf{n}}$). But the latter expression converges to $\int \chi(y) \ell(z) d\lambda(y, z)$ by definition of weak convergence, and this integral is equal in turn to $\int \chi(y) \psi(w) d\bar{\lambda}(y, z)$ by reversing the above argument. This completes the proof of continuity.

To complete the proof, it suffices to show that $\text{gr}(\Psi C)$ is the image of $\text{gr}(\hat{\Psi} C)$ under the projection that takes (x, λ) into $(x, \bar{\lambda})$. Suppose first that $(x, \lambda) \in \text{gr}(\hat{\Psi} C)$. Then the payoff from λ is

$$\begin{aligned} \int u_i(x, y, z) d\lambda(y, z) &= \int \left[\int u_i(x, y, z) d\nu(z|y) \right] d\mu(y) \\ &= \int u_i(x, y, \int z d\nu(z|y)) d\mu(y) \end{aligned}$$

(by the linearity of u_i)

$$= \int u_1(x, y, \bar{\nu}(\cdot | y)) d\bar{\mu}(y)$$

(by definition of $\bar{\nu}$). But $\bar{\nu}(\cdot | y) \in \text{co}[C(x, y)]$ because $\text{supp}[\nu(\cdot | y)] \subset C(x, y)$, and $\int u_1(x, y, \bar{\nu}(\cdot | y)) d\bar{\mu}(y)$ is by definition the payoff from $\bar{\lambda}$. So it is easy to see that $(x, \bar{\lambda}) \in \text{gr}(\Psi C)$.

Suppose now that $(x, \bar{\lambda}) \in \text{gr}(\Psi C)$. We need to find $(x, \lambda) \in \text{gr}(\hat{\Psi} C)$ of which $(x, \bar{\lambda})$ is the projection. To this end, let $\hat{C}(x, y, z)$ be the set of $\kappa \in \mathcal{PK}(Z)$ with expectation z and such that $\text{supp}[\kappa] \subset C(x, y)$. Then \hat{C} has a closed graph, and standard considerations show that $\hat{C}(x, y)$ is non-empty iff $z \in \text{co}[C(x, y)]$. Let $\hat{c}: \text{gr}(\text{co}[C]) \rightarrow \mathcal{PK}(Z)$ be a measurable selection from the restriction of \hat{C} to $\text{gr}(\text{co}[C])$, let $\nu(\cdot | y) = \hat{c}(x, y, \bar{\nu}(\cdot | y))$ for all $y \in A(x)$, and let $\nu(\cdot | y)$ be arbitrary for $y \notin A(x)$ (for example, one could set $\nu(\cdot | y) = \delta_{\bar{\nu}(\cdot | y)}$ for such y). Then it can be checked as above, using the linearity of u_1 , that the measure λ obtained by combining the marginal $\bar{\mu}$ with the transition probability ν lies in $\hat{\Psi} C(x)$. \square

Lemma 4.2 is the basic result justifying the backwards program. We conclude the present subsection by relating it to the work of Simon and Zame (1990).

The standard existence theorem for normal-form games can be expressed as follows. If each player's strategy set is a compact metric space, and if there is a continuous function associating a vector of payoffs with each vector of strategies, then there exists a vector of (possibly mixed) strategies that forms a Nash equilibrium. Simon and Zame (1990) extended this result by showing that if, instead, there is an upper semicontinuous correspondence associating a convex set of payoffs with each vector of strategies, then there exists a selection from this correspondence and a vector of (possibly mixed) strategies that forms a Nash equilibrium when the payoff function is given by this selection.

In order to explain how Lemma 4.2 extends the result of Simon and Zame, it is helpful to reinterpret their work slightly. Much as we do in the proof of Lemma 4.1, they consider a sequence of finite approximations to their basic game. Suppose these

approximations are indexed by N , suppose that the mixed strategies μ_{N_i} constitute a Nash equilibrium of game N for some selection from its payoff correspondence, and let the equilibrium payoffs of this equilibrium be π_{N_i} . Then their proof shows that, if the mixed strategies μ_i and payoffs π_i constitute a limit point of the strategies μ_{N_i} and payoffs π_{N_i} , then there exists a selection from the payoff correspondence of the original game such that the μ_i constitute a Nash equilibrium with equilibrium payoffs π_i when the payoff function is given by this selection. In other words, they demonstrate a limited form of upper semicontinuity for a particular continuous family of games.

Lemma 4.2 therefore extends the result of Simon and Zame in two ways.¹² The first is relatively minor: it shows that upper semicontinuity holds for a general continuous family of games. The second is more substantive. The result of Simon and Zame tells us about the behaviour of the equilibrium strategies and equilibrium payoffs, but it tells us very little about the selections from the payoff correspondence that enforce these strategies. Lemma 4.2, by contrast, concerns a comprehensive description of equilibrium. Such a description of equilibrium is not without intrinsic interest. But its main significance for the purposes of the present paper is that it allows us to avoid a measurable selection problem that would otherwise arise when we come to program forwards.¹³

¹² As it stands Lemma 4.2 is not a direct generalisation of the result of Simon and Zame. But it can easily be adapted to obtain one. To do so, view Z , not as a space of probability measures over a compact metric space, but more generally as a compact metrisable subset of a locally convex topological vector space. (This perspective is more general, and includes the possibility that Z is a compact subset of \mathbb{R}^I . If Z is a compact subset of \mathbb{R}^I , then one can define $u_i(x,y,z) = z_i$.) And define $\Psi C(x)$ in a more limited way, to consist of pairs $(\mu, \pi) \in \mathcal{PM}(Y) \times \mathbb{R}^I$, where μ is the distribution of a Nash equilibrium and π is its equilibrium payoff vector. (This more limited definition of ΨC is the price one must pay for the more general Z .) Then a proof almost identical to that of Lemma 4.2 shows that ΨC is a continuous projection of $\hat{\Psi C}$, and therefore upper semicontinuous.

¹³ It would be tempting to conduct the backwards program as follows. Let $\Pi: \text{gr}(A) \rightarrow \mathbb{R}^I$ be an upper-hemicontinuous convex-valued correspondence describing possible continuation payoffs. Then $\Psi\Pi: X \rightarrow \mathbb{R}^I$, the correspondence describing possible Nash equilibrium payoffs for the current period, is upper hemicontinuous; and so therefore is $\text{co}[\Psi\Pi]$, the continuation-payoff correspondence for the next iteration of the backwards program. This simple program is, however, difficult to reverse. For, in order to reverse it,

4.3 The Forwards Program

In this subsection we develop the analysis necessary for the forwards program. For this purpose we shall need a measurable selection c from ΨC . This function can be thought of as part of the output of previous iterations of the forwards program. It describes the equilibrium paths that must be followed from the current period onwards if the strategies calculated for earlier periods by the previous iterations of the forwards program are really to constitute part of an equilibrium.

Lemma 4.3 There exist measurable functions $\lambda: X \times [0,1] \rightarrow \mathcal{PM}(Y \times W)$ and $\nu: X \times Y \times [0,1] \rightarrow \mathcal{PM}(W)$ such that, for all $(x, \xi) \in X \times [0,1]$:

- (i) the marginal $\mu(\cdot | x, \xi)$ of $\lambda(\cdot | x, \xi)$ over Y is a product measure such that $\text{supp} [\mu(\cdot | x, \xi)] \subset A(x)$;
- (ii) the marginal $\mu_\emptyset(\cdot | x, \xi)$ of $\mu(\cdot | x, \xi)$ over Y_\emptyset is $f_\emptyset(\cdot | x)$;
- (iii) $\nu(\cdot | x, \cdot, \xi)$ is an r.c.p.d. of $\lambda(\cdot | x, \xi)$ over W such that $\nu(\cdot | x, y, \xi) \in \text{co}[C(x, y)]$ for all $y \in A(x)$;
- (iv) the marginals $\mu_i(\cdot | x, \xi)$ of $\mu(\cdot | x, \xi)$ over the Y_i constitute a Nash equilibrium when the continuation path following outcome y is $\nu(\cdot | x, y, \xi)$ for all $y \in A(x)$.

Moreover $\int \lambda(\cdot | x, \xi) d\xi = c(x)$ for all $x \in X$.

Lemma 4.3 captures the essence of the forwards program. It requires some explanation. First of all, conditions (i)–(iv) imply that $\lambda(\cdot | x, \xi) \in \Psi C(x)$. The condition $\int \lambda(\cdot | x, \xi) d\xi = c(x)$ therefore tells us that, if the equilibrium path is chosen from $\Psi C(x)$ on the basis of the current public signal $\xi \in [0,1]$, then the overall path will be precisely

one must be able to find, for any measurable selection π from $\Psi\Pi$, a measurable selection $\Phi\pi$ from Π such that the game with payoff function $\Phi\pi(x, \cdot)$ has a Nash equilibrium with payoff $\pi(x)$. This is not a standard measurable-selection problem. For what is actually required is that we select, for each x , a function $\Phi\pi(x, \cdot)$, and that these choices of functions be coordinated in x to obtain a single measurable function $\Phi\pi$. (A version of this non-standard problem did arise for Mertens and Parthasarathy (1987; Lemma 1 of Section 6, pp 41–42). In their case, however, the functions $\Phi\pi(x, \cdot)$ could be taken to be continuous.)

$c(x)$. (If Lebesgue measure is the marginal distribution over $[0,1]$, and if $\lambda(\cdot | x, \cdot)$ is the conditional distribution over $Y \times W$, then $\int \lambda(\cdot | x, \xi) d\xi$ is the marginal distribution over $Y \times W$.) Secondly, the lemma goes beyond the simple assertion that there exists $\lambda: X \times [0,1] \rightarrow \mathcal{PM}(Y \times W)$ such that $\lambda(\cdot | x, \xi) \in \Psi C(x)$ for all (x, ξ) . For, while this assertion would automatically imply the existence of a ν that was measurable in y for each (x, ξ) , the proposition delivers a ν that is measurable in (x, y, ξ) jointly. Thirdly, the $\mu_i: X \times [0,1] \rightarrow \mathcal{PM}(Y_i)$ are precisely the strategies of the players for the current period. Fourthly, the restriction Φc of ν to $\text{gr}(A) \times [0,1]$ describes the equilibrium paths that must be followed from the next period onwards. It serves as the input for the next iteration of the forwards program.

Proof Most of the considerations necessary for the proof have already arisen in the proof, of Lemmas 4.1 and 4.2. The present proof will therefore be brief.

The first step is to find a measurable $d: X \rightarrow \mathcal{PM}(\mathcal{PM}(Y \times W))$ such that, for all $x \in X$, $\text{supp}(d(x)) \subset \Psi C(x)$ and the expectation of $d(x)$ is $c(x)$. This can be done using an argument very similar to one used in the proof of Lemma 4.2. Next, we need a measurable mapping $\lambda: X \times [0,1] \rightarrow \mathcal{PM}(Y \times W)$ such that, when $[0,1]$ is given Lebesgue measure, the distribution of the random variable $\lambda(\cdot | x, \cdot): [0,1] \rightarrow \mathcal{PM}(Y \times W)$ over $\mathcal{PM}(Y \times W)$ is precisely $d(x)$ for all x . Such a mapping can be obtained by applying a ‘measurable’ representation theorem, e.g. Gihman and Skorohod (1979; Lemma 1.2, p 9). The third step is to find a measurable mapping $\tilde{\nu}: X \times Y \times [0,1] \rightarrow \mathcal{PM}(W)$ such that $\tilde{\nu}(\cdot | x, \cdot, \xi)$ is an r.c.p.d. of $\lambda(\cdot | x, \xi)$ for all $(x, \xi) \in X \times [0,1]$. In other words, we need a ‘measurable’ version of the standard result which asserts the existence, for each (x, ξ) , of an r.c.p.d. $\tilde{\nu}(\cdot | x, \cdot, \xi)$ of $\lambda(\cdot | x, \xi)$. Since this is, in effect, the central step of the forwards program, and since we have not been able to find precisely the result we need in the literature, we sketch a proof of this.

Let $\mu(\cdot | x, \xi)$ be the marginal of $\lambda(\cdot | x, \xi)$ over Y . Let $\{B_n | 1 \leq n < \infty\}$ be a sequence of sets that generates the Borel σ -algebra on Y . For each N , let \mathcal{F}_N be the

σ -algebra generated by $\{B_n | 1 \leq n \leq N\}$. And let $\bar{\nu}$ be a fixed element of $\mathcal{PM}(W)$.

Since \mathcal{F}_N is finite we can give an explicit formula for an r.c.p.d. of $\lambda(\cdot | x, \xi)$ given the σ -algebra \mathcal{F}_N on Y . For each $y \in Y$ and each measurable $\tilde{W} \subset W$ we set

$$\nu_N(\tilde{W} | x, y, \xi) = \frac{\lambda(\tilde{Y} \times \tilde{W} | x, \xi)}{\mu(\tilde{Y} | x, \xi)}$$

if \tilde{Y} is the atom of \mathcal{F}_N containing y and $\mu(\tilde{Y} | x, \xi) > 0$, and

$$\nu_N(\tilde{W} | x, y, \xi) = \bar{\nu}(\tilde{W})$$

otherwise. It is immediate from these formulae that $\nu_N(\cdot | x, y, \xi)$ is jointly measurable in (x, y, ξ) , and therefore that the convergence set K of ν_N is measurable. Set $\tilde{\nu} = \lim_{N \rightarrow \infty} \nu_N$ on K and $\tilde{\nu} = \bar{\nu}$ outside K .

Certainly $\tilde{\nu}$ is measurable. Also, since $\nu_N(\cdot | x, \cdot, \xi)$ is effectively a version of the conditional expectation of $\lambda(\cdot | x, \xi)$ given \mathcal{F}_N , the martingale convergence theorem implies that $\lim_{N \rightarrow \infty} \nu_N(\cdot | x, \cdot, \xi)$ exists $\mu(\cdot | x, \xi)$ -a.s., and that $\tilde{\nu}(\cdot | x, \cdot, \xi)$ is a version of what is effectively the conditional expectation of $\lambda(\cdot | x, \xi)$ given the Borel σ -algebra on Y . That is, $\tilde{\nu}(\cdot | x, \cdot, \xi)$ is an r.c.p.d. of $\lambda(\cdot | x, \xi)$. This completes the third step.

There are two difficulties with $\tilde{\nu}$ as it stands: it may not be the case that $\tilde{\nu}(\cdot | x, \cdot, \xi) \in \text{co}[C(x, y)]$ for all $y \in A(x)$; and the marginals $\mu(\cdot | x, \xi)$ of $\lambda(\cdot | x, \xi)$ need not constitute a Nash equilibrium when the continuation path following outcome y is $\tilde{\nu}(\cdot | x, y, \xi)$. The fourth and final step constructs a ν that does have these properties from $\tilde{\nu}$. To obtain the first property, let q be any measurable selection from $\text{co}[C]$, and redefine $\tilde{\nu}(\cdot | x, y, \xi)$ to be $q(x, y)$ wherever it does not lie in $\text{co}[C(x, y)]$ for some $y \in A(x)$. To obtain the second: for each i , let q_i be the function defined in the proof of Lemma 4.1; let $\pi_i(x, \xi) = \int u_i(x, y, \tilde{\nu}(\cdot | x, y, \xi)) d\mu(y | x, \xi)$; and let $\tilde{Y}_i(x, \xi)$ be the set of a_i

$\in A_i(x)$ such that $\int u_i(x,y|a_i, \tilde{\nu}(\cdot|x,y,\xi))d\mu(y|x,\xi) > \pi_i(x,\xi)$. Then we may set $\nu(\cdot|x,y,\xi) = q_i(x,y)$ if $y_j \in Y_j \setminus \tilde{Y}_j(x,\xi)$ for all $j \neq i$ but $y_i \in \tilde{Y}_i(x,\xi)$, and $\nu(\cdot|x,y,\xi) = \tilde{\nu}(\cdot|x,y,\xi)$ otherwise. Standard considerations show that the resulting ν is measurable, and the fact that $\lambda(\cdot|x,\xi) \in C(x)$ implies that $\nu(\cdot|x,\cdot,\xi)$ is still an r.c.p.d. of $\lambda(\cdot|x,\xi)$. \square

4.4 Completion of the proof

Before proceeding, we need to define the equilibrium correspondence for a family of games. In making this definition we encounter a minor technicality. We have defined an equilibrium as a strategy profile that is in Nash equilibrium in every subgame. But a strategy profile specifies players' behaviour for the entire family of games that they might face. Hence an equilibrium will generate a whole family of paths, one for each game $x^0 \in X^0$. We therefore define $E_1(x^0)$, the set of equilibrium paths of game x^0 , to be the set of paths λ such that λ is the equilibrium path generated in game x^0 by some equilibrium. Similarly, reinterpreting our basic family of games as a family parameterised by $x^{t-1} \in X^{t-1}$, we can define correspondences E_t for $t \geq 2$.

Theorem 4.1 Suppose that (A1)–(A3) hold. Then E_t is an upper semicontinuous map from X^{t-1} into $\mathcal{H}(\mathcal{PM}(\times_{s=t}^{\infty} Y_s))$.

The main interest of Theorem 4.1 derives from the fact that it implies existence (see Section 5). However upper semicontinuity is itself of interest: it is reassuring to know that our concept of equilibrium possesses this standard regularity property.

Proof The proof proceeds in three steps. Let $B_t(x^{t-1})$ be the set of paths beginning in period t that are consistent with Nature's strategy. The first part of the proof shows that B_t is an upper semicontinuous map from X^{t-1} into $\mathcal{H}(\mathcal{PM}(\times_{s=t}^{\infty} Y_s))$. In the second, a sequence of correspondences $\{C_t | 1 \leq t < \infty\}$ is constructed, each of which is upper

semicontinuous. In the third it is shown that $E_t = \text{co}[C_t]$ for all t .

For notational convenience, we show that B_1 is upper semicontinuous. To this end, let $B_1^\dagger(x^0)$ be the set of marginals over $\times_{s=1}^t Y_s$ of measures in $B_1(x^0)$. Then a backwards programming argument similar to, but simpler than, that developed in Lemmas 4.1 and 4.2 shows that B_1^\dagger is an upper semicontinuous map from X^0 into $\mathcal{H}(\mathcal{PM}(\times_{s=1}^t Y_s))$. But a sequence of measures $\{\lambda_n\} \subset \mathcal{PM}(\times_{t=1}^\infty Y_t)$ converges iff its marginals over $\mathcal{PM}(\times_{s=1}^t Y_s)$ converge for all t . So B_1 , too, is upper semicontinuous. This completes the first step.

Turning to the second step, let $T \geq 1$ be given. For each $t > T$, let $C_t^T = B_t$. And, for each $1 \leq t \leq T$, let C_t^T be defined inductively by the formula $C_t^T = \Psi_t(C_{t+1}^T)$. (In other words, $C_t^T(x^{t-1})$ is the set of equilibrium paths that are obtained for period t onwards when the continuation paths can be chosen from $\text{co}[C_{t+1}^T]$.) Then the backwards program ensures that each C_t^T is upper semicontinuous.

Now define $C_t = \cap_{T=1}^\infty C_t^T$ for each $t \geq 1$. It is obvious that $C_t^T \subset C_t^S$ whenever $T \geq S$. Hence C_t too is upper semicontinuous. It is also the case that $\Psi_t(C_{t+1}) \subset C_t$. To see this, let $\lambda \in [\Psi_t(C_{t+1})](x^{t-1})$. By definition, λ can be enforced by continuation paths from C_{t+1} . But paths in C_{t+1} belong a fortiori to C_{t+1}^T . Hence $\lambda \in [\Psi_t(C_{t+1}^T)](x^{t-1})$ for all T . Since $C_t = \cap_{T=1}^\infty C_t^T$, the required conclusion follows. Finally, $C_t \subset \Psi_t(C_{t+1})$. To see this note that the correspondences C_{t+1}^T for $1 \leq T < \infty$ together with C_{t+1} for $T = \infty$ form a continuous family of games with parameters x^{t-1} and T . And Lemma 4.2 shows that the equilibrium correspondence of such a family is upper semicontinuous. It follows that any $\lambda \in C_t(x^{t-1})$, which is certainly a limit point of points in $[\Psi_t(C_{t+1}^T)](x^{t-1})$, also lies in $[\Psi_t(C_{t+1})](x^{t-1})$.

Overall, then, we obtain a sequence of upper semicontinuous correspondences $\{C_t \mid 1 \leq t < \infty\}$ such that $C_t = \Psi_t(C_{t+1})$ for all t . It remains only to show that $E_t = \text{co}[C_t]$ for all t . We do this for $t = 1$ only. The other cases are analogous.

We show first that $\text{co}[C_1] \subset E_1$. It suffices to show that any measurable selection c_1 from $\text{co}[C_1]$ is also a selection from E_1 . So let c_1 be any such selection. Then

repeated application of the forwards program starting with c_1 yields, for all $t \geq 1$, strategies f_{ti} and a measurable selection c_{t+1} from $\text{co}[C_{t+1}]$ such that: (i) for all $(x^{t-1}, \xi^t) \in H^{t-1}$, the $f_{ti}(x^{t-1}, \xi^t)$ constitute a Nash equilibrium in period t when the continuation paths are given by $c_{t+1}(x^{t-1}, \cdot, \xi^t)$ and (ii) the paths obtained from period t onward when the f_{ti} are employed in period t and the continuation paths are given by c_{t+1} are precisely those required by c_t . It follows from (i) and a standard dynamic programming argument that the strategy profile f constitutes an equilibrium. And it follows from (ii) that the equilibrium path in game x^0 is precisely $c_1(x^0)$. This completes the proof that $\text{co}[C_1] \subset E_1$.

We now show that $E_1 \subset \text{co}[C_1]$. To this end, let $f = \langle f_{ti} | t \geq 1, i \in \mathcal{I} \rangle$ be any subgame-perfect equilibrium in extended strategies. For each $t \geq 1$, let $c_t: H^{t-1} \rightarrow \mathcal{P}\mathcal{K}(X_{s=t}^{\text{w}} Y_s)$ describe the paths followed from period t onwards when the strategy profile f is employed. The mapping c_t inherits measurability from f . Let $T \geq 1$ be given. Now certainly c_{T+1} is a selection from $\text{co}[C_{T+1}^T]$. Also, suppose that c_{t+1} is a selection from $\text{co}[C_{t+1}^T]$ for some $1 \leq t \leq T$. Since f is an equilibrium, the f_{ti} specify Nash equilibria for period t when the continuation paths are given by c_{t+1} . It follows at once that c_t is a selection from $\text{co}[C_t^T]$. Proceeding inductively we therefore conclude that c_1 is a selection from $\text{co}[C_1^T]$. Moreover, allowing T to vary, we obtain that c_1 is a selection from $\text{co}[C_1]$. This completes the proof. \square

5 Existence with and without Public Signals

This section provides a systematic discussion of the question of existence of equilibrium in the dynamic model introduced in Section 3. It begins with a recapitulation of the main result for this case, namely that, when public signals are allowed, any family of games satisfying the standing assumptions (A1)–(A3) possesses an equilibrium. However, the section also addresses two further questions. Under what circumstances can existence be obtained without the use of public signals? And is the extension of the concept of subgame–perfect equilibrium obtained by allowing the observation of public signals sufficiently rich?

5.1 Existence with Public Signals

The following result is an immediate corollary of Theorem 4.1:

Theorem 5.1 Suppose that (A1)–(A3) hold. Then there exists a subgame–perfect equilibrium in extended strategies.

Proof Pick any $x^0 \in X^0$. By Theorem 4.1, $E_1(x^0)$ is non–empty. Pick any $\lambda \in E_1(x^0)$. Then, by definition of $E_1(x^0)$, there exists an equilibrium of the family of games that generates equilibrium path λ in game x^0 . In particular, there exists an equilibrium!

□

The potential drawback of this result is that, in order to obtain the existence equilibrium in what is intended to be a purely non–cooperative situation, it introduces an extraneous element into the model, namely public signals. It is therefore of some interest to discover circumstances under which equilibrium exists even without this element.

5.2 Existence without public signals

There are at least three classes of game in which public signals are not needed for existence: games of perfect information, finite-action games, and zero-sum games. The results for such games are summarised in Theorems 5.2, 5.3, and 5.4 respectively.

Theorem 5.2 Suppose that (A1)–(A3) hold, and that the family of games has perfect information. Then a subgame-perfect equilibrium exists. \square

Proof Oddly enough, Theorem 5.2 is not an immediate corollary of Theorem 5.1.

(Intuitively it would seem obvious that, when only one player acts at a time, public signals do not allow players to achieve anything they could not achieve using mixed strategies.

But there is a minor complication: players observe not only the current signal, but also past signals.) However, an equilibrium can be constructed using the correspondences C_t used in the proof of Theorem 4.1. We begin with any measurable selection c_1 from C_1 , and define $f_{1i}(x^0)$ to be the marginal of $c_1(x^0)$ over Y_{1i} . Next, let c_2 be a measurable selection from $\text{co}[C_2]$ that enforces c_1 . Since the game is a game of perfect information, $\text{co}[C_2] = C_2$. Hence c_2 is also a measurable selection from C_2 . So we may define $f_{2i}(x^1)$ to be the marginal of $c_2(x^1)$ over Y_{1i} . And so on. \square

Indeed, one can even show that there exists a subgame-perfect equilibrium in which players employ pure actions only. We do not demonstrate this here, since the argument does not appear to be a simple corollary of the construction that we have employed.¹⁴ Note that the proof of Theorem 5.2 shows incidentally that allowing public signals in a game of perfect information does not enlarge the set of equilibrium paths.

¹⁴ See Harris (1985) for a treatment of the case of perfect information. The argument used there can easily be adapted to the present context. Alternatively, one can work directly with the apparatus used here, taking care to purify the selections as one programs forward.

Let us say that a family of games is of finite action if X^0 is finite, and if, for all $t \geq 1$ and all $x^{t-1} \in X^{t-1}$, $A_t(x^{t-1})$ is a finite set.

Theorem 5.3 Suppose that (A1)–(A3) hold, and that the family of games is of finite action. Then a subgame–perfect equilibrium exists. Moreover, for each $x^0 \in X^0$, the set of equilibrium paths is compact.

Proof Theorem 5.3 is more or less well known, and there is more than one way of proving it. (See Fudenberg and Levine (1983) for a different proof.) So we merely note that it can be proved by a simple adaptation of the methods of this paper. The basic idea can be expressed in terms of the notation of Section 4.2 as follows. Instead of defining $\Psi C(x)$ to be the set of equilibrium paths that can be obtained by choosing continuation paths from $\text{co}[C]$, one defines it to be the set of equilibrium paths that can be obtained by choosing continuation paths from C . Because X is finite, and because the $A_i(x)$ are finite for all $x \in X$, this new definition still yields an upper semicontinuous correspondence.

It should be noted that both finiteness assumptions are needed. Indeed, suppose that X is finite but that one or more of the $A_i(x)$ are infinite for some x . Then it can happen that there is a sequence of equilibria of game x , each of which involves a public signal generated endogenously by players' randomisation over their action sets $A_i(x)$, and that this sequence converges to a limiting equilibrium which does not generate any such signal. Suppose, on the other hand, that the $A_i(x)$ are finite for all $x \in X$ but that X is infinite. Then it can happen that there is a convergent sequence of games in X , and that there is at least one player for whom at least one player for whom at least two actions coalesce in the limit. If this player randomises over his two actions, then he generates a public signal which disappears in the limit.

It is precisely in order to compensate for the potential disappearance of such endogenously generated public signals that the exogenous public signals are used in our basic existence theorem. \square

To complete our discussion of the existence of subgame—perfect equilibrium, we consider the case of a family of zero—sum games. (A family of games is of zero sum if \mathcal{G} contains precisely two elements, and if $\sum_{i \in \mathcal{I}} u_i \equiv 0$.)

Theorem 5.4 Suppose that (A1)—(A3) hold, and that the family of games is of zero sum. Then a subgame—perfect equilibrium exists.

In order to understand this result, it is helpful to recall some facts about one—period zero—sum games. It is well known that each players' set of optimal strategies is convex in such a game. But the set of equilibria, regarded as the set of probability distributions over outcomes, need not be. Hence, even in a zero—sum game, the set of equilibria will be enlarged if players are allowed to observe a public signal. The set of equilibrium payoffs, on the other hand, will not. For the public signal merely allows the two players to co—ordinate on a choice of Nash equilibrium for the game, and all Nash equilibria have the same equilibrium payoff (namely $(v, -v)$, where v is the value of the game).

Proof Note first that our proof of the existence of equilibrium shows, in particular, that the extended family of games has a Nash equilibrium. It follows immediately from standard considerations that this family also has a value. That is, there exists a measurable mapping $v_1 : X^0 \rightarrow \mathbb{R}$ such that $v_1(x^0)$ is the value of extended game x^0 . (Indeed, since $v_1(x^0)$ is simply player 1's payoff from any path in $E_1(x^0)$, and since E_1 is upper semicontinuous, v_1 is actually continuous. We do not need this fact, however.) Similarly, there exist mappings $v_t : X^{t-1} \rightarrow \mathbb{R}$ such that $v_t(x^{t-1})$ is the value of the extended game beginning in period t .

With these considerations in mind, we can construct a subgame—perfect equilibrium of the original family of games as follows. The construction follows the construction of an equilibrium of the extended game, but takes care to avoid the need for a

public signal. The first step is to find a measurable selection \tilde{c}_1 from C_1 . Since \tilde{c}_1 is a measurable selection from C_1 rather than $\text{co}[C_1]$, its marginals over Y_1 are already product measures. So we may define strategies $f_{1i}: X^0 \rightarrow \mathcal{P}\mathcal{M}(Y_{1i})$ by setting $f_{1i}(x^0)$ to be the marginal of $\tilde{c}_1(x^0)$ over Y_{1i} . Similarly, \tilde{c}_1 generates a measurable selection c_2 from $\text{co}[C_2]$ that enforces the f_{1i} . But all $\lambda \in \text{co}[C_2(x^1)]$ generate the same continuation payoff, namely $v_2(x^1)$. In particular, if \tilde{c}_2 is a measurable selection from C_2 then \tilde{c}_2 generates the same continuation payoffs as c_2 . Hence the f_{1i} can be enforced by \tilde{c}_2 just as well as by c_2 . (The continuation paths specified by \tilde{c}_2 do not necessarily constitute an r.c.p.d. of \tilde{c}_1 , of course.) Iterating this argument we obtain strategies $f_{ti}: X^{t-1} \rightarrow \mathcal{P}\mathcal{M}(Y_{ti})$ for all $t \geq 2$ as well. \square

5.3 Approximation by finite-action games

Suppose that we are given a single game (i.e. a family of games for which X^0 is a singleton) satisfying (A1)–(A3). Then a finite-action approximation to this game is, roughly speaking, a new family of games indexed by $n \in \mathbb{N} \cup \{\infty\}$ such that: the game corresponding to $n = \infty$ is the original game; for each $n \in \mathbb{N}$, the game corresponding to n is of finite action; and taken as a whole the new family of games satisfies (A1)–(A3). (Here (A1) ensures that players' action sets in the original game are continuously approximated; (A2) ensures that Nature's behaviour is continuously approximated; and (A3) ensures that active players' payoffs are continuously approximated.)

Finite-action approximations to a single game satisfying (A1)–(A3) certainly exist. Indeed, we can even construct a finite-action approximation to a family of games satisfying (A1)–(A3). To do this, we need the following ingredients. First, for each $t \geq 1$ and each $i \in \mathcal{I}$, let $\{y_{ti}^n | n \in \mathbb{N}\}$ be dense in Y_{ti} ; and let $\{y_0^n | n \in \mathbb{N}\}$ be dense in $Y_0 = X^0$. Secondly, for each $n \in \mathbb{N}$, let $Y_0^n = \{y_0^m | 1 \leq m \leq n\}$. Thirdly, for each $t \geq 1, i \in \mathcal{I}$ and $n \in \mathbb{N}$, let $\tilde{\alpha}_{ti}^n: X^{t-1} \rightarrow Y_{ti}$ be a continuous function such that, for all $x^{t-1} \in X^{t-1}$: $\tilde{\alpha}_{ti}^n(x^{t-1})$ belongs to $A_{ti}(x^{t-1})$, and $\tilde{\alpha}_{ti}^n(x^{t-1})$ is as close to y_{ti}^n as any other point in $A_{ti}(x^{t-1})$. (Such a function exists because $A_{ti}: X^{t-1} \rightarrow \mathcal{K}(Y_{ti})$ is continuous. In

Nature's case the function is simply constant with value $y_{t\emptyset}^n$.) And let $\alpha_{ti}^n: X^{t-1} \rightarrow \mathcal{K}(Y_{ti})$ be defined by $\alpha_{ti}^n(x^{t-1}) = \{\tilde{\alpha}_{ti}^m(x^{t-1}) \mid 1 \leq m \leq n\}$ for all $x^{t-1} \in X^{t-1}$.

Fourthly, for each $t \geq 1$ and $n \in \mathbb{N}$, let $\varphi_{t\emptyset}^n: X^{t-1} \rightarrow \mathcal{PM}(Y_{t\emptyset})$ be a continuous function such that, for all $x^{t-1} \in X^{t-1}$: $\text{supp}[\varphi_{t\emptyset}^n(x^{t-1})] \subset \{y_{t\emptyset}^m \mid 1 \leq m \leq n\}$; and $\varphi_{t\emptyset}^n(x^{t-1})$ is as close to $f_{t\emptyset}(x^{t-1})$ as another measure in $\mathcal{PM}(Y_{t\emptyset})$ whose support is contained in $\{y_{t\emptyset}^m \mid 1 \leq m \leq n\}$. Finally: let $Y_0^{\omega} = X_0$; for all $t \geq 1$, $i \in \mathcal{I}$ and $x^{t-1} \in X^{t-1}$, let $\alpha_{ti}^{\omega}(x^{t-1}) = A_{ti}(x^{t-1})$; and for all $t \geq 1$ and $x^{t-1} \in X^{t-1}$ let $\varphi_{t\emptyset}^{\omega}(x^{t-1}) = f_{t\emptyset}(x^{t-1})$.

In terms of these ingredients we can define a larger family of games, which can be interpreted as a sequence of families of finite-action games approximating the original family, as follows. First, let \hat{X}^0 be the union over $n \in \mathbb{N} \cup \{\omega\}$ of the sets $\{n\} \times Y_0^n$. Secondly, suppose that \hat{X}^{t-1} has been defined for some $t \geq 1$. Then: for all $i \in \mathcal{I}$, \hat{A}_{ti} is the restriction of $\alpha_{ti}^{\cdot}(\cdot)$ to \hat{X}^{t-1} ; $\hat{A}_t = \times_{i \in \mathcal{I}} \hat{A}_{ti}$; and $\hat{X}_t = \text{gr}(\hat{A}_t)$. Thirdly, for all $t \geq 1$, $\hat{f}_{t\emptyset}$ is the restriction of $\varphi_{t\emptyset}^{\cdot}(\cdot)$ to \hat{X}^{t-1} . Finally, payoff functions \hat{u}_i are defined in the obvious way.

Now suppose that we are given a finite-action approximation to a single game. Since this approximation is simply a family of games parameterised by $n \in \mathbb{N} \cup \{\omega\}$ and satisfying (A1)-(A3), its equilibrium correspondence is upper semicontinuous in n . Hence, since any subgame-perfect equilibrium path of any game $n \in \mathbb{N}$ is, in particular, an equilibrium path of that game, it follows that any limit point of subgame-perfect equilibrium paths of the finite-action games is an equilibrium path of the original game. And our solution concept (namely subgame-perfect equilibrium in extended strategies) includes everything that can be obtained by taking finite-action approximations to a game (or, more precisely, everything that can be obtained by taking continuous finite-action approximations to a game satisfying (A1)-(A3)).

5.4 Approximation by ϵ -equilibria

One can also show that any limit point of subgame-perfect ϵ -equilibrium paths of a given single game is an equilibrium path of the game, by combining the techniques of this

paper with those of Börgers (1989). Indeed, suppose that we are given a family of games satisfying (A1)-(A3). For each $t \geq 1$, $x^{t-1} \in X^{t-1}$ and $\epsilon \geq 0$, let $\hat{E}_t(\epsilon, x^{t-1})$ be the set of equilibrium paths of subgame x^{t-1} of a family of games identical to the original game except that: (i) each player $i \in \mathcal{N}$ is replaced by a sequence $\{t_i | t \geq 1\}$ of agents; (ii) each agent chooses actions that are at or within ϵ of being optimal in each subgame. Then it is easy to adapt the methods of Section 4 to show that \hat{E}_t is upper semicontinuous. But this implies the required result. For any subgame-perfect ϵ -equilibrium path of the original game is also an ϵ -equilibrium path; any ϵ -equilibrium path is an ϵ -equilibrium path of the game in which each active player is replaced by a sequence of agents; and the 0-equilibrium paths of the game in which each active player is replaced by a sequence of agents coincide with the equilibrium paths of the original game.

7 Conclusion

The results of this paper concerning the existence of subgame-perfect equilibrium in general dynamic games create a small dilemma. If one adopts the point of view that the observation of public signals violates the spirit of non-cooperative game theory, then the interpretation of the counterexample of Section 2 is presumably that games with a continuum of actions are somehow intrinsically pathological. After all, that example seems wholly non-pathological in all other respects. If, on the other hand, one thinks that games with a continuum of actions are an indispensable part of non-cooperative game theory, then presumably one must accept that games must be extended to allow the observation of public signals.


Faced with this dilemma it is tempting to dismiss games with a continuum of actions as pathological. There are, however, at least two reasons why this reaction may be too simplistic. First, it can be argued that a solution concept should be robust to small errors in the data of the games to which it applies, in the sense that one should not rule out equilibria that occur in games arbitrarily close to the game under consideration.¹⁵ (In other words, the equilibrium correspondence should be upper semicontinuous.) Provided that such errors are confined to payoffs, the solution concept of subgame-perfect equilibrium satisfies this requirement when applied to finite-action games. If, however, such errors may be present in the extensive form itself, then public signals must be introduced if upper semicontinuity is to be satisfied. Secondly, it is not entirely clear that public signals really do violate the spirit of non-cooperative game theory. After all, players must in any case coordinate on a specific Nash equilibrium to play; and public signals merely increase the extent to which they can coordinate, by allowing them to coordinate on a random Nash equilibrium.

¹⁵ See Fudenberg, Kreps and Levine (1988) for example.

References

- Börgers, T (1989): "Perfect Equilibrium Histories of Finite- and Infinite-Horizon Games", Journal of Economic Theory 47, 218–227.
- Chakrabarti, S K (1988): "On the Existence of Equilibria in a Class of Discrete-Time Dynamic Games with Imperfect Information", mimeo, Department of Economics, Indiana University.
- Duffie, D, J Geanakoplos, A Mas-Colell & A McLennan (1989): "Stationary Markov Equilibria", Research Paper 1079, Graduate School of Business, Stanford University.
- Forges, F (1986): "An Approach to Communication Equilibria", Econometrica 54, 1375–1385.
- Fudenberg, D, D Kreps & D Levine (1988): "On the Robustness of Equilibrium Refinements", Journal of Economic Theory 44, 354–380.
- Fudenberg, D, & D Levine (1983): "Subgame-Perfect Equilibria of Finite- and Infinite-Horizon Games", Journal of Economic Theory 31, 251–268.
- Fudenberg, D, and J Tirole (1985): "Pre-Emption and Rent Equalization in the Adoption of a New Technology", Review of Economic Studies 52, 383–401.
- Gihman, I I, and A V Skorohod (1979): Controlled Stochastic Processes, Springer-Verlag.
- Harris, C (1985a): "Existence and Characterisation of Perfect Equilibrium in Games of Perfect Information", Econometrica 53, 613–628.
- Harris, C (1985b): "A Characterisation of the Perfect Equilibria of Infinite-Horizon Games", Journal of Economic Theory 37, 99–125.
- Harris, C (1990): "The Existence of Subgame-Perfect Equilibrium with and without Markov Strategies: a Case for Extensive-Form Correlation", Discussion Paper 53, Nuffield College, Oxford.
- Hellwig, M, & W Leininger (1986): "A Note on the Relation between Discrete and Continuous Games of Perfect Information", Discussion Paper A-92, University of Bonn.
- Hellwig, M, & W Leininger (1987): "On the Existence of Subgame-Perfect Equilibrium in Infinite-Action Games of Perfect Information", Journal of Economic Theory 43, 55–75.
- Mertens, J-F, & T Parthasarathy (1987): "Equilibria for Discounted Stochastic Games", Research Paper 8750, CORE, Université Catholique de Louvain.
- Parthasarathy, K R (1967): Probability Measures on Metric Spaces, Academic Press.
- Simon, L K, & W R Zame (1990): "Discontinuous Games and Endogenous Sharing Rules", Econometrica 58, 861–872.

5000 472




Date Due

| Date Due | | |
|----------|--|--|
| | | |

Lib-26-67

MIT LIBRARIES DUPL



3 9080 00672155 6

